

EXHIBIT 18

EXHIBIT 6



Intel® Technology Journal

The Original 45nm Intel Core™ Microarchitecture

Intel Technology Journal Q3'08 (Volume 12, Issue 3) focuses on Intel® Processors Based on the Original 45nm Intel Core™ Microarchitecture: The First Tick in Intel's new Architecture and Silicon "Tick-Tock" Cadence

Original 45nm Intel® Core™ 2 Processor Performance

**The Technical Challenges of Transitioning
Intel® PRO/Wireless Solutions to a Half-Mini Card**

**Power Management Enhancements
in the 45nm Intel® Core™ Microarchitecture**

Greater Mobility Through Lower Power

**Improvements in the Intel® Core™ 2
Penryn Processor Family Architecture
and Microarchitecture**

Power Improvements on 2008 Desktop Platforms

**Mobility Thin and Small Form-Factor Packaging
for Intel® Processors Based on Original 45nm
Intel Core™ Microarchitecture**

The First Six-Core Intel® Xeon™ Microprocessor

More information, including current and past issues of Intel Technology Journal, can be found at:

<http://developer.intel.com/technology/itj/index.htm>



Intel® Technology Journal

The Original 45nm Intel Core™ Microarchitecture

Articles

Preface	iii
Foreword	v
Technical Reviewers	vii
Original 45nm Intel® Core™ 2 Processor Performance	157
Power Management Enhancements in the 45nm Intel® Core™ Microarchitecture	169
Improvements in the Intel® Core™ 2 Penryn Processor Family Architecture and Microarchitecture	179
Mobility Thin and Small Form-Factor Packaging for Intel® Processors Based on Original 45nm Intel Core™ Microarchitecture	193
The Technical Challenges of Transitioning Intel® PRO/Wireless Solutions to a Half-Mini Card	199
Greater Mobility Through Lower Power	211
Power Improvements on 2008 Desktop Platforms	219
The First Six-Core Intel® Xeon® Microprocessor	229

The First Six-Core Intel[®] Xeon[®] Microprocessor

Shankar Sawant, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Ravishankar Kuppaswamy, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Anantha Kinnal, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Kuldeep Simha, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Ravindra Saraf, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Pradeep Kaushik, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Srikanth Balasubramaniam, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Narayanan Natarajan, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Gautam Doshi, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Sambit Sahu, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Mysore Sriram, Enterprise Microprocessor Group, DEG, Bangalore, Intel Corporation
 Jeffrey Gilbert, Xeon Architect, DEG/DAP/MAP, Hillsboro, OR., Intel Corporation

Index words: multi-core, front-side-bus, energy-efficient, 3-level cache hierarchy, high-K metal gate, 45nm

ABSTRACT

This paper describes the next-generation Intel[®] Xeon[®] microprocessor designed for a broad range of highly power-efficient servers, codename Dunnington. The Dunnington processor has six cores (three core-pairs) integrated with large, dense, on-chip caches, and it delivers the dramatic power efficiency of Intel's 45nm high-K metal gate process and the Intel Core™ 2 microarchitecture to server platforms. This processor implements a high bandwidth-dedicated interface from each of the three core pairs to the last-level cache (LLC) for effective use of the inclusive LLC. With high functional integration, large cache size, and 1.9 billion transistors, the processor's moderate server-class die size of 503 mm² is achieved by optimizing the floor plan and physical design. This six-core product is designed to be a plug-in refresh for server platforms, codename Caneland, using Intel's DHSI-based quad-core processor, codename Tigerton. The Dunnington processor will be offered in multiple options with core counts of four or six, LLC sizes of 12 or 16 MB, several core frequencies, and Thermal Design Power (TDP) limits of 50, 65, 90, and 130 W. This processor will be the first part to employ core recovery techniques for reducing product cost. Compared to the Tigerton processor, it provides an average performance improvement of more than 25 percent. The benefits of the 45nm hi-K process and the Penryn family of processors' base is seen in the doubling of Performance/Watt

and in the low TDP limits that permit six-core compute capability in the blade-server form factor.

INTRODUCTION

The high-performance expandable server processor market segment has witnessed ever-increasing demands for throughput performance and energy efficiency. To meet these demands we focused on intelligent integration of multiple cores for power-efficient parallel computing to deliver increased performance. The key high-level design requirements for the next-generation Intel[®] Xeon[®] microprocessor, codename Dunnington, Intel's latest offering in this segment, were a drop-in replacement compatibility with its predecessor, Intel's Dedicated High Speed Interconnect (DHSI)-based quad-core processor, codename Tigerton, on the Caneland platform; a 30-percent boost in performance over its predecessor; and operation in the range of 50–130 W power envelopes. In addition, maintaining compatibility between different pools of machines and stacks of software is one of the major problems in data-center and server management. Hence, enhancing the virtualization support for load-balancing across computing pools was a critical feature requirement for the Dunnington processor. This is the first Intel six-core Xeon processor, integrating a three-level, on-chip cache hierarchy and a fast DHSI system interface slated for introduction in the second half of 2008. The increased number of cores

meant a higher memory bandwidth from the DHSI interface of the Caneland platform, thereby requiring improved cache organization over its predecessor. The requirement for integrating six cores and such a large Last-Level Cache (LLC) had profound implications on the physical and electrical design of the Dunnington processor. Significant among these were die size constraints, achieving bin-split targets for the different market segments, meeting reliability constraints due to the 1.9 billion on-die transistors, dealing with process variations across the large die, and meeting thermal envelope limits. In this paper we describe each of these challenges and the solutions developed by our team to successfully bring the product to market ahead of schedule.

ARCHITECTURE FEATURES

Growing utilization of the Front-Side Bus (FSB) bandwidth required multiple architectural solutions in order to keep the soaring bandwidth-latency curve under control. Two solutions, introduced in the Caneland platform, include DHSI and snoop filters. However, to meet the required performance on the Caneland platform, the traditional multi-chip-package-based solutions for multi-core processors provided little opportunity to effectively address this challenge. The monolithic hex-core solution adopted by the Dunnington processor design allows it to tackle the bandwidth-latency challenge using an efficient cache hierarchy, a high-bandwidth on-die interface, and innovative solutions to reduce snoop traffic, thus providing a compelling Intel Xeon processor product for Q3 2008.

As illustrated in Figure 1, the Dunnington design consists of three Penryn processor family CMP core-pairs that are integrated with the LLC and the Caching Bridge Controller (CBC) using a point-to-point protocol called Simple Direct Interface (SDI). SDI provides improved latencies and bandwidth for LLC and cross-core data accesses as compared to similar accesses in its predecessor (which take place via the FSB). The CBC that stands between the Penryn cores and the LLC has three different roles to play in the Dunnington design: (a) it is a coherence and conflict resolution engine for data access from three core-pairs and the external snoops; (b) it is a cache controller for LLC and core-to-core data transfers; and (c) it is an FSB controller for the Dunnington processor.

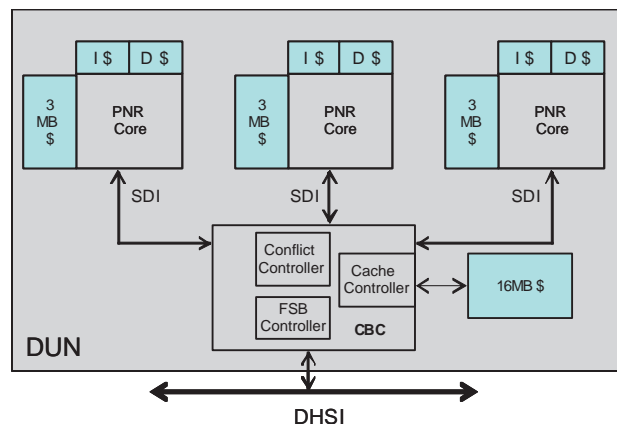


Figure 1: Next-generation Intel® Xeon® processor architecture

CACHE ORGANIZATION

The Dunnington processor has three levels of caches: 32 KB of data and 32 KBs of instruction cache in each Penryn Core (First-Level Cache or FLC), 3 MB of non-inclusive Mid-Level Cache (MLC) for each CMP core-pair, and 16 MB of inclusive LLC. Inclusivity of the LLC is maintained using core valid bits per Penryn core-pair.

The Dunnington processor implements the joint MESI states of ES, MI, and MS. (With joint MESI states, the first letter reflects the caching privileges of the LLC, while the second letter indicates the caching privileges granted to a lower level cache—in the case of Dunnington, FLC and MLC in each Penryn core-pair). The LLC holds all lines at an MESI state with equivalent or greater privilege than the FLC and the MLC have for their associated lines. The objective of joint MESI states is to optimize response time to external snoops without materially penalizing internal cache privilege transfer.

The MLC and LLC are safeguarded by Intel Cache Safe Technology. The MLC utilizes the Single Bit Fix (SBF) mechanism, while the LLC includes a cache line disable facility. These features address the cache reliability requirements of the server market segment.

COMPATIBILITY FEATURES

The processor's FSB interface is designed to be compatible with Tigerton's FSB and to operate at the highest FSB frequency (1066MT/s) in use on multi-processor platforms during the product's lifetime.

Active-way management (AWM) is another key compatibility and performance feature on the Caneland platform. AWM improves the effectiveness of the

Caneland platform's snoop filter and the efficiency of the FSB by sending way-hints to the Clarksboro chipset. These way-hints align snoop filter victimization with processor cache victimization. The near-flawless tracking of cache line allocation avoids invalidations of active lines in the LLC and the resultant increased effective memory latency.

The Dunnington processor continues to support all power-management features from its predecessor. These include P-states (P0, P1), C-states (C1E, CC3), and T-states (TM1, TM2).

VIRTUALIZATION ENHANCED

Most data centers add computational capacity over time. New generations of processors and platforms invariably offer a better price per performance and generally offer reduced operating cost per platform. Additionally, they offer reduced operating cost per unit performance. The cost efficiencies of next-generation server products means most data-center populations will have more than one processor family deployed and may also have different platforms.

There can be compatibility issues across these population factions. Typically, a next-generation processor supports incremental features to its prior-generation processor. These changes mean it becomes harder to maintain a common software stack between these factions or to move computation dynamically between platforms in the different factions. V-Migration (also known as V-Motion) is a feature that addresses these software compatibility issues by providing the ability to virtually demote a next-generation processor to function as a current-generation processor from an instruction set standpoint. It must be noted that performance and power requirements remain equivalent to the next-generation processor, providing the end-user the advantages of the next-generation technology. With this feature implemented on the Dunnington processor, the Caneland platform enables a seamless virtual machine transfer from previous generations to the Dunnington processor.

In order to provide aggressive performance per watt and to fit in various power/performance envelopes of the server segment, multiple SKUs are available by using variants of core counts, cache sizes, and core frequencies. The CBC architecture is designed to effectively adapt itself to these variations.

With these architectural innovations, the Dunnington processor delivers as much as an overall 30-percent gain on existing workloads for the same power envelopes when compared to its quad-core predeces-

sor. Hence, it provides a significant boost to power-performance efficiency for the Caneland platform.

PROCESS TECHNOLOGY

One of the key innovations in the new 45nm process technology is the high-k + metal gate transistor, which is one of the biggest changes in transistor technology since the introduction of the polysilicon gate MOS transistor in the late 1960s. The new 45nm process technology offers about a $2 \times$ improvement in transistor density, approximately 20 percent of an improvement in transistor switching speed, or more than a $5 \times$ reduction in the source-drain leakage. It also provides at least a $10 \times$ reduction in gate oxide leakage power and more than a 30-percent reduction in transistor switching power [1].

GLOBAL ELECTRICAL CHALLENGES

The Dunnington product is a large die that integrates six cores, a 16-MB LLC, and the Uncore logic. We found these key electrical challenges:

- Achieving the required bin-splits in the different market segments on core frequency.
- Meeting the Uncore frequency targets on the FSB-digital logic.
- Meeting Vmax reliability constraints associated with the huge number of transistors on the die.
- Meeting Vmin constraints imposed by the cache memory cells and register files on the core.
- Meeting the Thermal Design Power (TDP) for the product.

Systematic and random variations of the manufacturing process parameters introduced design constraints such as pre-silicon frequency targets on different design domains, error-correction/redundancy requirements on the cache, and compensation mechanisms on the global clocks. Identifying the optimal process targeting to achieve the SKU stackup of core count, core frequency, cache size, and TDP power requirements is a significant challenge.

In these subsequent sections we discuss how we met these challenges in the different design domains.

PHYSICAL DESIGN

The Dunnington processor die integrates three dual-cores from the Penryn family of processors, 9 MB of MLC, and 16 MB of LLC in just over 500 mm^2 . The floor plan of the chip is shown in Figure 2.

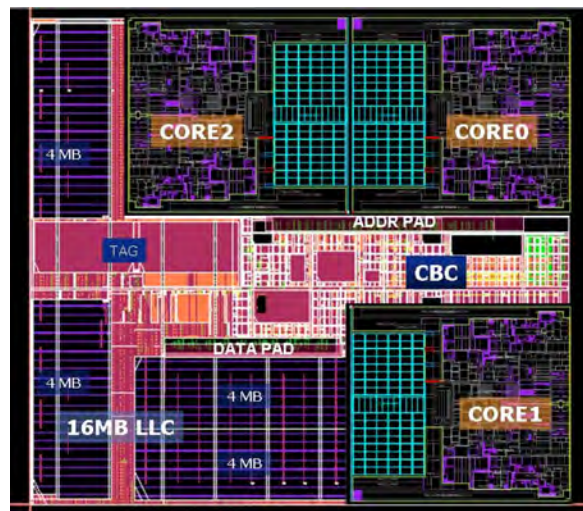


Figure 2: Six-core processor die photo

The Dunnington design presented unique physical design challenges. The first of these was fitting all of these components into a single die, subject to constraints on the maximum die size, module orientation constraints, and tight timing requirements. Die size limits prevented all three blocks from the Penryn family of processors from being placed in one line, so we had to place them in two rows as shown. I/O blocks, traditionally placed at the edges of the die, had to be accommodated in the center of the die to enable IO operation at the target FSB speed. This required us to carefully optimize the floor plan as there was heavy wiring congestion in this area. Early block size estimates and routing analyses were critical to establish the feasibility of the floor plan. Clever modular design of the LLC floor plan enabled the cache team to fit 16 MB of cache into the irregular shape that remained after the cores and IOs had been placed. This saved a significant amount of custom layout effort for the project.

Another major challenge was repeater insertion on global signals. Due to the size of the chip, virtually all global signals required repeaters to meet signal integrity constraints. We had to insert over 50,000 repeaters into the LLC and CBC, in addition to the 55,000 repeaters that already existed inside each of the three Penryn modules. We followed a “virtual repeater” methodology in order to avoid having to code these repeaters into Run-Time Library (RTL) to enable faster timing convergence. We used a fully automated virtual repeater insertion tool along with a correct-by-construction flow to convert over 700 virtual repeater “stations” into physical layout blocks after timing convergence was achieved. Having a guaranteed flow to make the virtual repeaters “real” without requiring manual layout cleanup enabled the

design to stay in virtual repeater mode until very late in the design cycle. This flexibility allowed rapid timing convergence progress.

We implemented most of the Uncore logic using automated synthesis and place-and-route tools, as part of the overall project focus on schedule. Early engagement with our EDA tool vendor to enable 45nm design rule support was crucial and was done in close partnership with another internal 45nm design project. The block design teams did careful analyses of the optimal block size; the merged small blocks and split overly large blocks to get the best timing and layout convergence behavior.

Module layout IP reuse from other internal projects was a key theme in the physical design of the product. Apart from the cores that were reused from the Penryn family of processors, we also reused other hard IPs such as Phased Lock Loops (PLLs), cache sub-arrays, and analog IO cells from other products. While this approach saved a significant amount of custom layout design and validation effort, this reuse of IP did create complexities of its own, such as layout grid mismatches, different tool/flow environments, and contrasting design methods, all of which needed creative integration solutions. A chip with 1.9 billion transistors is bound to stress layout completion and verification tools to the limit. Meeting stringent volume-manufacturing design rules for large die, such as metal and via density and minimum and maximum spacing rules, required sophisticated automation tools and flows. The design team used customized layout verification tools so that they were more efficient in pinpointing layout issues in the huge database. At time of final tape-in, the size of the layout database was over 89 GB, which was very stressful on the compute servers and network infrastructure.

CACHE DESIGN

The Dunnington processor has a three-level cache architecture as described earlier. Each 3-MB MLC is logically organized as 4 K lines by 12 ways. Data are transferred in cache line quantities (64 bytes) across a 256-bit bus with a maximum bandwidth of 32 bytes per clk. MLC accesses are pipelined to permit a new request every two clocks and have a latency of 9 cycles from a request to data return at MLC interface. Additional details relating to the physical implementation of the MLC are reported in [2].

The Dunnington processor has 16 MB of unified LLC organized as 16 ways, 16-K sets, and 64-byte cache lines. Physically the 16-MB data cache is organized in 4-MB blocks, each containing 4-K sets and 16 ways as shown in Figure 3. The lower-cache-size SKUs are

supported by reducing the number of ways to 12 or 8. Each 4-MB block is further made up of 1-MB sub-blocks containing 16 data banks. A data bank comprises 4 sub-arrays and a mid logic. Each sub-array is divided into two 256-wordline halves with 296 bitlines. Intel's 45nm Ultra Low Voltage (ULV) cell is chosen as the memory cell for its small size and robust low-CC performance. All data and control signals are staged through the mid logic en route to the sub-arrays. The LLC data design is fully synchronous, that is, access to the nearest array has the same latency as the access to the spatially farthest array. The irregular shape of the LLC on the die results in a physically different implementation of the horizontal 4-MB blocks when compared to the vertical 4-MB blocks. The differences include changes in the routing topology and data flow, muxing schemes, and physical repeater placements. The data cache is protected by Single Error Correct, Double Error Detect (SEC-DED) Error Correction Scheme (ECC) performed inline.




Figure 3: 16 MB LLC, 9 MB MLC, and 1.5 MB tag on the six core die

The 1.5-MB tag cache is physically implemented as two sections each containing 8-K sets. Structurally, each section comprises 16 ways, with each way further consisting of two sub-arrays. State and core valid bits are also stored in the tag sub-arrays. The tag arrays are also protected by an inline SEC-DED ECC scheme.

Power-saving features form a key aspect of the LLC design. Aggressive clock-gating schemes reduce dynamic power. The datapath latching stages are controlled using gated clocks to eliminate unnecessary clock transitions. Fine-grained sleep transistor implementation reduces leakage power. During an LLC access, only one sixteenth of a 1-MB slice is active, reducing both active and leakage power. The remain-

ing portion of the sub-array remains in a sleep mode in which the power to the SRAM cells is dropped below the nominal voltage. Sufficient redundancy is implemented in the design to improve large cache yields.

CLOCK DOMAINS AND DISTRIBUTION

The Dunnington processor has three primary clock domains. The first is the high-frequency core clock domain (GCLK) that supports the Penryn family of processors' core and its associated L2 cache. The second is the half-core-frequency clock domain (SCLK) that supports most of the Uncore logic and the LLC. And the third domain is the quad-pumped FSB clock (ZCLK) that serves the FSB pads and the common clock signals. The GCLK frequency is an integer multiple (8, 9, 10, or 11) of the input clock (xxCLK), the SCLK is always one-half the GCLK frequency, and the ZCLK is always four times the input clock frequency. A fixed GCLK:SCLK ratio of two was chosen for faster time-to-market by reducing design and validation time. Figure 4 illustrates the clock system architecture for the Dunnington processor. Each Penryn family of processors' core has two embedded PLLs. The first of these is the IO PLL that receives the xxCLK and synthesizes the 4X frequency DCLK. The IO PLL also generates a reference clock for the second PLL, called the core PLL. The core PLL generates the high-frequency core clock required by the core logic. This cascade of the IO PLL driving the core PLL is replicated for Uncore. Thus, the processor has a total of eight PLLs, of which six are embedded in the three Penryn family of processors' cores; the remaining two PLLs support the Uncore. Multiple S-macros (SCLK-Macros) are placed at the end of the Uncore GCLK distribution to generate a half-frequency SCLK for Uncore.

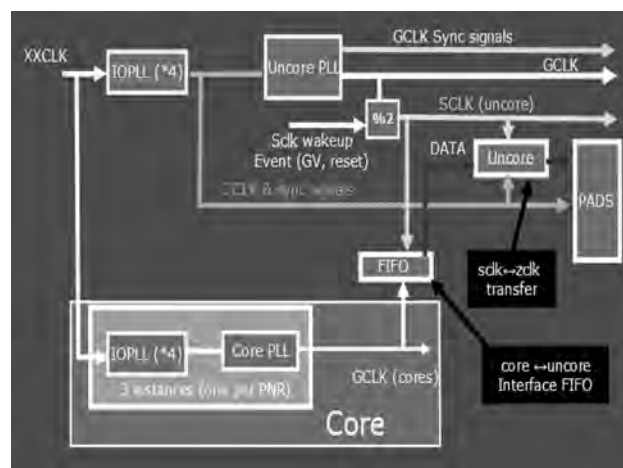


Figure 4: The processor clock architecture

Given the large die size and points-of-divergences that are four PLLs apart, the skew at the Core-Uncore boundary can be large. The Core-Uncore communication is enabled through a rate matched (GCLK to SCLK) FIFO. The pointer separation is programmable to multiples of GCLK cycles allowing for optimization post silicon to achieve higher throughput. A similar protocol is implemented at the Uncore pad boundary, although it is not designed to be optimized post-silicon due to a shorter point of divergence between the Uncore IO and the Core PLLs.

A single set of C4 bumps receives the differential input clock, xxCLK. This is necessitated by package routing resources that are constrained due to the location of the central FSB pads. The xxCLK is routed on-die, in a balanced tree structure, as is the reference clock to the four IO PLLs. Fuse programmable delays, added to the reference clock path to the four IO PLLs, allow for tuning of any systemic skew between them post-silicon. Although the FSB clock inside the core is mostly redundant, it is preserved for the reference clock generation for its core PLL. Its distribution network is preserved to maintain the integrity of the feedback clock path to the IO PLL. Similarly, the GCLK distribution in the core is preserved as well to maintain the integrity of the distributions inside the cores. Figure 5 shows the Uncore clock distribution spines and the distribution topology.

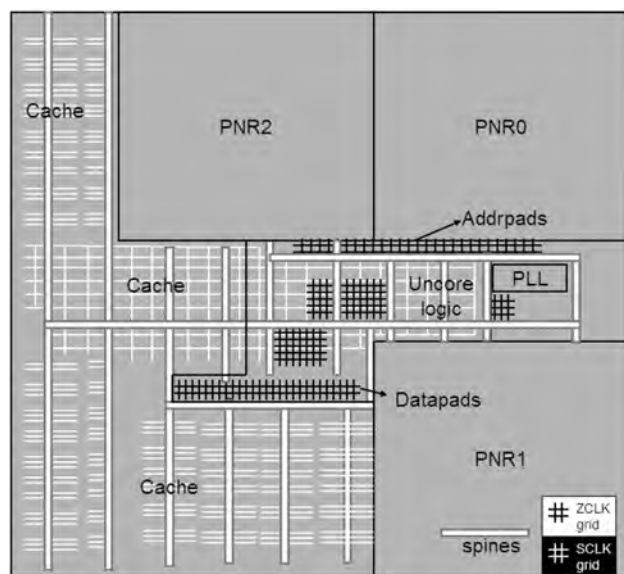


Figure 5: Uncore clock distribution

A new GCLK distribution is created for the Uncore. GCLK is distributed to the entire Uncore via 18 vertical spines, with one horizontal spine acting as the backbone feeding them. At the root of each vertical spine, programmable delays are inserted to offset skew

mismatches post-silicon. SCLK is generated in the Uncore at the tail end of distribution by combining the GCLK with the latency matched SclkSync signal. Tight skew control is achieved by creating a SCLK grid over the key logic areas, and power optimization is achieved by depopulating the SCLK clock grid lines. For areas such as the data arrays of the LLC, which are more skew tolerant, a point-to-point distribution is implemented to achieve additional power savings.

IO AND PACKAGING

The Dunnington processor IO is a point-to-point DHSI design running at 1067 MT/sec and is socket compatible with its predecessor, the Tigerton processor, on the Caneland platform. Unlike conventional microprocessor designs that place the IO pads at the edges of the die, often separating one set of IO buffers from another by the entire length/width of the die, in the Dunnington processor the address and data buffers are placed at the center of the die. Locating the IOs close to the Uncore logic and to each other allows mitigation of several internal timing-critical paths, and this is key to hitting 1066MT/s. In addition to relocating the pads, other architectural changes to improve the timing between the processor and chipset have been made. These include new features such as the ADS-enabled inbound address path and a clocking scheme that decouples the platform timing constraints from the processor's bus ratio.

The basic architecture and analog design collateral for the IO buffers are reused from the Penryn family of processors. All circuit and architectural changes are made in a controlled fashion to significantly accelerate design convergence.

While the movement of the pads to the center of the die ensured that the latency and, in turn, the performance goals are met, it certainly introduces a slew of challenges for the package design. The IO buffers typically use a bump pitch that is different from the rest of the processor. Therefore, locating the IO buffers in the center of the die complicates the bump pattern significantly. Figure 6 shows the multiple bump patterns that result from the new IO location. The alternating bump patterns cause an uneven epoxy underfill flow, resulting in voids during packaging. Several experiments were conducted using an Assembly Test Vehicle to ensure that this problem was understood and corrected prior to actual fabrication.




Figure 6: Bump transitions on the die

A second challenge related to the new IO location is the design of the package routes. Figure 7 shows the routing solution designed to allow signals to escape from the center of the die without interfering with signal integrity.




Figure 7: Package routing scheme

POWER CONSUMPTION

With the added number of cores in a process generation, the overall power dissipation of the processor increases proportionally. This affects both the leakage power and the dynamic power components. The Dunnington processor, however, was required to fit multiple market segments, namely rack, blade and ultra-dense segments, which have TDP requirements of approximately 130, 90, and 60 W, respectively. Initial analysis of the microarchitecture indicated that the thermal envelope would be violated if steps were not taken to address the exposure. Power dissipation can be reduced by lowering the voltage with a commensurate reduction in frequency. However, the voltage cannot be reduced below the VCCmin target, because this may affect functional robustness. The VCCmin target is determined primarily by the overall spread of statistical variation [2] of all transistors in a

power plane. Hence, operating voltage reduction to reduce power is an option only for VCCmin. The Dunnington processor was designed to be socket compatible with a quad-core Tigerton processor on the Caneland platform. The specifications of the Caneland platform mandate one digital power plane only, and hence the entire digital logic including the cores, caches, and Uncore had to reside on this power plane. This meant over 1.9 billion transistors are operated on a single power plane, which affects the VCCmin of the product considerably. Hence, further measures besides voltage lowering were required to fit the six cores, the large cache, and the Uncore into a thermal envelope of 130 W. Figure 8 shows the initial breakup of leakage and dynamic power for the various functional areas.




Figure 8: TDP breakup with no power reduction

To reduce power without violating VCCmin targets, the microprocessor design had to be re-targeted toward a low-leakage version of Intel's 45nm hi-K metal-gate process technology [1]. The low-leakage version of the process reduced leakage by a factor of three beyond the significant advancements that the Intel 45nm process brought with it. However, with the low-leakage process option, the transistor delays increase by approximately 13 percent. The low-leakage process does not affect the VCCmin of the product, and hence, for a marginal reduction in frequency, a significant power reduction was achieved. Figure 9 shows the leakage and dynamic power distributions after applying the low-leakage process to the product. As the figure shows, more power is allocated to the dynamic component, enabling a higher frequency of operation while simultaneously improving the power performance of the microprocessor.




Figure 9: TDP breakup after re-targeting design to low-leakage process

To counter the frequency loss due to low-leakage process targeting, guard bands were incorporated into the timing analysis of the design. Further, clock and other critical aspects of the design that affect functional robustness were simulated using Monte Carlo analysis to ensure the low-leakage process targeting did not affect functionality or yield, and guaranteed robust performance.

POWER DELIVERY

As mentioned in the previous section, the Dunnington processor has a single power plane for all digital transistors. Because 1.9 billion transistors are operated off one power plane, the VCCmin target increases for the power plane. However, due to the increased transistor count, the maximum voltage (VCCmax) is reduced. This is affected due to reliability concerns, also known as Gox reliability or gate oxide reliability. The reduction of VCCmax with the simultaneous increase in VCCmin tightened the operating voltage range for the product. Hence, the power delivery for such a narrow voltage operating range had to be carefully designed so as not to cause functional or reliability issues.

The Dunnington power-delivery solution was designed around three power planes. The digital power plane or VccCore is the largest power plane and operates between 0.85 and 1.1 v. The FSB IO as well as fuse and thermal circuits are serviced by the VTT power plane, which operates at 1.1 v. The Vanalog power plane is used only within the various PLLs in the chip. Figure 10 shows the three power planes used on the product.




Figure 10: Power domains represented on the die (VccCore in blue, VTT in red, Vanalog in green)

The designers were confronted with three primary power-delivery challenges. Firstly, with a 130-W TDP envelope, coupled with a low-leakage process, the di/dt noise on the power rails could cause large first droops that could have impacted frequency as well as functionality. Detailed microarchitecture studies were done by the architecture and circuitry team to determine the precise nature of the current ramp. Further, to reduce the voltage droops, accurate modeling of the silicon die was done as the load to the di/dt was extracted, and simulations were done to optimize the on-die as well as package capacitance components.

Secondly, due to Caneland platform constraints and placement of the FSB IO pads in the center of the die, the VTT power delivery to critical IO and fuse analog structures was affected. To counter this, the package traces delivering VTT to the center of the die were strengthened. Further, an additional row of bumps was allocated in the center of the die for VTT power supply to the pads and fuse regions.

Finally, due to the large die multi-core nature of this chip, long busses had to be routed between the various cores, the Uncore, and the caches, necessitating large repeater stations. These repeater stations had very high power density, causing a large IR drop. Because some of the repeater stations reside adjacent to the cache arrays, the increased IR drop would have caused VCCmin failures. The power grids over repeater stations were therefore strengthened through improved metal layer coverage for both power and ground connections. A further, high amount of explicit on-die decoupling capacitance was placed within the repeater stations to mitigate IR drop.

DESIGN FOR TEST AND DESIGN FOR MANUFACTURABILITY

Most DfX features implemented on this product are inherited from the Penryn family of processors' core and are extended for the Uncore. Traditional features include Scan-Out, FRC; and DCM, BIST, LYA for arrays, I/O Loopback test, and pattern generator for FSB, etc. Enhancements done by the core, such as fuse programmability for analog blocks, parallel testing of large arrays (data/tag), and modes extending granularity in LCPs, were retained on the product.

The most significant DfX features added were site selection and site symmetry. For efficient re-use of High Volume Manufacturing (HVM) content on each of the three Penryn family of processors' core sites on the die and test-time reduction, both of these features are very important. With site-symmetric design, functional patterns generated for one site were identical cycle-by-cycle to those at the other two sites. Just by changing the fuse pattern for selecting a core site, the same patterns could be re-run on the other two sites on the HVM tester. Because of the symmetry of the design, we could also produce identical signatures for FRC tests, thereby giving automatic coverage to the second core on the same site, thus reducing total test time. Core selection was also necessary to support a 4-core SKU.

LOCK-STEP USE MODEL

The Dunnington processor design leveraged the FRC architecture implemented in a single site. For HVM coverage, the use model was to have internal-FRC enabled between the two cores within a single site, core0 within the site acting as the master and core1 within the site acting as a slave. FRC across the three sites (0, 1, and 2) was not supported, since the motivation was to reuse legacy Penryn HVM content to achieve maximum coverage for each site. Likewise, external-FRC, with injected data traffic and in-between SITE0/1/2, is not supported due to limitations in the availability of routing space and spare in-die interconnects between the Uncore and each of the Penryn family of processors' sites. For the Dunnington design, SIGMODE was enabled for the sites only and not for the Uncore logic.

SIGMODE DESIGN IMPLEMENTATION

Site0, Site1, and Site2 contain approximately 31,000 scanout nodes. Internal to a site, these nodes are distributed into seven sub-chains named BUS, BLS, L2, FRC0, FRC1, CORE0, and CORE1. These sub-chains begin and end at the individual site's TAP

controller and can be multiplexed in a variety of configurations.

For SIGMODE, these have been configured into two chains feeding into the Linear Feedback Shift Registers (LFSR) that compress the signature data captured from the scanout nodes. To read the final signature, each site is selected using a TAPSEL feature, and the signature is serially shifted out one-by-one from the LFSRs (Figures 11 and 12).

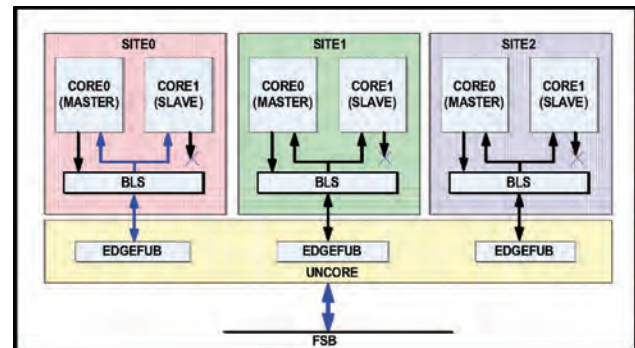


Figure 11: FRC or lockstep enabled on single site, two cores operating in master and slave mode

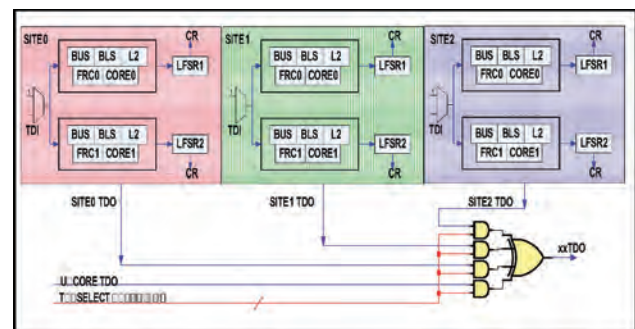


Figure 12: Topology of SIGMODE chain connectivity on the die for three sites

CONCLUSION

The Dunnington processor team integrated six cores on a single die with a 25-MB cache to achieve a 30-percent increase in performance over its predecessor on the same platform. This product has virtualization features that are critical on servers for datacenter applications. Performance per watt efficiency was accomplished with a low-leakage variant of the 45nm process technology to enable SKUs for high-performance, rack, and ultra-dense segments. Modular design and reuse of IP/methodologies during the entire design and validation cycles enabled a successful execution that beat the time-to-market window.

ACKNOWLEDGEMENTS

We thank all of the architects, designers, and validators who collaborated in the creation of this product.

REFERENCES

- [1] Mistry K. *et al.* A 45nm Logic Technology with High-k + Metal Gate Transistors, Strained Silicon, 9 Cu Interconnect Layers, 193nm Dry Patterning, and 100% Pb-free Packaging. International Electron Devices Meeting, December 2007, pp. 247–250.
- [2] Varghese George *et al.* Penryn: 45nm Next Generation Intel Core 2 Processor. IEEE Association, November 2007.
- [3] Keln Kuhn *et al.* Managing Process Variation in Intel's 45nm CMOS Technology. Intel Technology Journal, Vol. 12, No. 2, 2008.

AUTHORS' BIOGRAPHIES

Ravi Kuppuswamy is the Project Manager for the next-generation Intel® Xeon® microprocessor, codename Dunnington, in the Enterprise Microprocessor Group in Intel's Digital Enterprise Group. In this role, he is responsible for all the post A0 design and development activities to get this server processor into high-volume manufacturing. Ravi joined Intel in 1996 and has held several technical and management positions across five generations of Intel's process lead-vehicle microprocessor programs. Ravi earned his Master's degree in Electrical Engineering from Arizona State University. He also has a Master's degree in Chemistry and a Bachelor's degree in Electrical Engineering from Birla Institute of Technology & Science, India. Ravi's e-mail is ravishankar.kuppuswamy at intel.com.

Kuldeep S. Simha is a Design Manager in Intel's Digital Enterprise Group, specializing in circuit design and global electricals on high-speed microprocessors. Kuldeep received a B.Tech degree in Electronics & Communication from the National Institute of Technology, Karnataka, India in 1993 and an M.S. degree in Computer Engineering from the University of Cincinnati in 1998. Prior to joining Intel in 2003, he worked in the area of telematics in CDOT and microprocessor design at Hewlett Packard. He has worked on a large number of product areas including register-file and cache design, global clocking design, and power/frequency analysis. Kuldeep's e-mail is kuldeep.s.simha at intel.com.

Shankar Sawant is an Engineering Manager in Intel's Enterprise Microprocessor Group. Shankar has been involved across four generations of Intel's process

lead-vehicle microprocessor programs in various technical and managerial roles. He received his Bachelors degree from the Indian Institute of Technology-Madras, India, and a Ph.D degree from North Carolina State University, Raleigh, in Electrical Engineering. Shankar's email is shankar.r.sawant at intel.com.

Anantha Kinnal is a Design Engineering Manager in the Enterprise Microprocessor Group at Intel India. He has been with Intel for over five years now and has worked on server processor designs. Previously he worked at Nexgen Microsystems on Nx-586/587 and then at AMD on K6 and K8 processor designs. His areas of interest include pre/post silicon validation, emulation, DFT/DFM, and platform architecture definition. His email is Anantha.kinnal at intel.com.

Mysore Sriram is a Principal Engineer with the Design and Technology Solutions Group. During his 15-year career with Intel, he has worked on the development of several internal EDA tools in the physical design domain, and he has worked on physical integration methodology and execution for several processor design projects. His research interests are in the areas of optimization algorithms, interconnect analysis and design, and place-and-route. His email is mysore.sriram at intel.com.

Pradeep Kaushik received a B.Tech degree in Electronics and Communication from the Delhi Institute of Technology in 1995. Prior to joining Intel, he worked in the areas of caches and IOs with SUN Microsystems. Since 2004, he has been with Intel. His technical interests are in the field of clocking and cache design. Currently, he leads the clocking team in the EMG server product team in India. His email is pradeep.kaushik at intel.com.

Ravi Saraf is a Design Engineer with Intel's Enterprise Microprocessor Group. He received his Bachelors degree from the College of Engineering-Pune, India, and a Master's degree from the Indian Institute of Technology, Bombay, India. He has been with Intel for six years and was involved in microarchitecture definition and the development of server processor designs. Ravi's research interests are in computer and platform architecture. His email is ravindra.p.saraf at intel.com.

Srikanth Balasubramanian joined Intel in 2003 and has been involved in the electrical design of server microprocessors. He has worked on various parts of electrical design of processors ranging from clock circuits, datapath design, power delivery, and bin splits. Recently for the next-generation Intel® Xeon® microprocessor, codename Dunnington, he was

responsible for global circuit methodology, power delivery, and fuse designs. Prior to joining Intel, he worked in the design of digital signal processors. His interests are in the field of electrical robustness of circuits, low skew clock delivery, low power design, and high-efficiency power delivery. Srikanth has a master's degree from IIT Madras.

Jeffrey Gilbert joined Intel in 1997 as the Willamette System Validation architect. After a brief time in the DPG Platform Architecture group, he joined the Oregon Design Center. Then in the newly formed Xeon Architecture group, he was the architect for two generations of server processors. Jeff received 2007 Intel Achievement Awards for his work on the Tulsa microprocessor and the MCP methodology used for Intel's initial quad-core products. This is Jeff's second time at Intel, the first being 1983–1987 when he worked on debug tools for the 80386 and P7 microprocessors. He was a software and hardware contractor in the interregnum. His primary interest is power-efficient platform performance for the server market segments. Jeff has an MSEE from Stanford University. He can be reached at jeffrey.d.gilbert@intel.com.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Intel's trademarks may be used publicly with permission only from Intel. Fair use of Intel's trademarks in advertising and promotion of Intel products requires proper acknowledgement.

Any codenames featured in this document are used internally within Intel to identify products that are in development and not yet publicly announced for release. For ease of reference, some codenames have been used in this document for products that have already been released. Customers, licensees, and other third parties are not authorized by Intel to use codenames in advertising, promotion or marketing of any product or services and any such use of Intel's internal codenames is at the sole risk of the user.

*Other names and brands may be claimed as the property of others.

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

Intel Corporation uses the Palm OS® Ready mark under license from Palm, Inc.

LEED—Leadership in Energy & Environmental Design (LEED®)

Copyright © 2008 Intel Corporation. All rights reserved.

This publication was downloaded from <http://www.intel.com>.

Additional legal notices at: <http://www.intel.com/sites/corporate/tradmarx.htm>.